

(19)

Europäisches Patentamt

European Patent Office

Office européen des brevets



(11)

EP 0 886 403 A1

(12)

EUROPEAN PATENT APPLICATION

(43) Date of publication:
23.12.1998 Bulletin 1998/52

(51) Int. Cl.⁶: H04L 12/56, H04L 29/06

(21) Application number: 97401433.4

(22) Date of filing: 20.06.1997

(84) Designated Contracting States:
AT BE CH DE DK ES FI FR GB GR IE IT LI LU MC
NL PT SE
Designated Extension States:
AL LT LV RO SI

(72) Inventor: Michiel, Herman
3078 Everberg (BE)

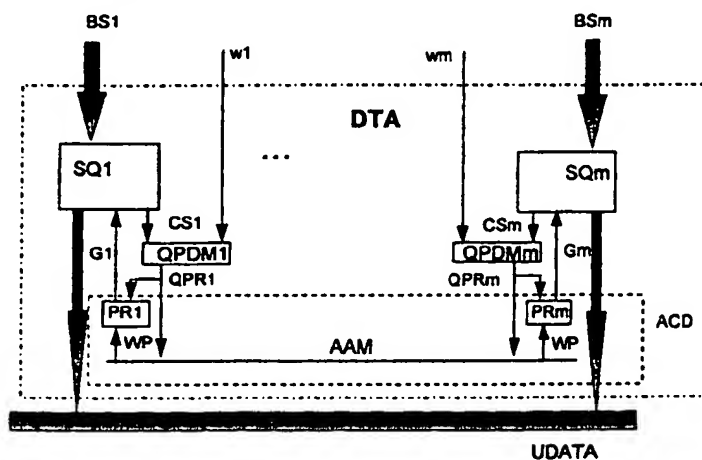
(74) Representative: Narmon, Gisèle
Industrial Property Department,
Alcatel Bell N.V.,
Francis Wellesplein 1
2018 Antwerpen (BE)

(71) Applicant:
ALCATEL ALSTHOM COMPAGNIE GENERALE
D'ELECTRICITE
75008 Paris (FR)

(54) Method and arrangement for prioritised data transmission of packets

(57) A method and arrangement for prioritised transmission of packets from a plurality of storage queues, each interfacing one corresponding input bitstream to one common transmission medium, is disclosed, whereby the respective queue priorities, on the basis of which access to the common transmission medium is determined, are itself dependent upon

respective weights associated to the corresponding input bitstreams. This dependency on these respective weights relates to a respective amount of successive values of these respective queue priorities, to be selected from an arbitrary interval and sorted in a pre-determined sequence.



EP 0 886 403 A1

BEST AVAILABLE COPY

Description

The present invention relates to a method and an arrangement for prioritised data transmission of packets as is further described in the preamble of claims 1 and 7.

Such an arrangement for performing said method is already known in the art, e.g. from the published European patent application *Data transmission system and interface module and priority generating means included therein*, publication number 0584405 A1. Therein, a data transmission system including a plurality of interface modules, each interfacing a plurality of input bitstreams with a common transmission medium, is described, whereby each interface module includes storage means as well as a packet priority determining means for determining a packet priority for each of the received packets. Within one interface module, the packet priorities are used in a priority generator for further determining an interface module priority, on the basis of which access to the transmission medium is determined by means of processing units included in each module and by means of an access bus included in the data transmission arrangement. The combination of these processing units and the access bus thus corresponds to the access control device as described in claim 7. In the special case, where a packet priority determining means for generating the packet priority, is attributed to one single storage queue, and where each interface module merely interfaces one input bitstream with the common transmission medium, the prior art arrangement corresponds to the arrangement as described in the preamble of claim 7 for realising the method described in the preamble of claim 1. Indeed, in this special case, the packet priority determining means and the priority generator of the prior art document together correspond to the queue priority determining means, described in the preamble of claim 7, whereby both the packet priority and the interface module priority correspond to the queue priority as described in the preamble of claims 1 and 7.

The prior art method and arrangement are suited for determining a packet priority, and in the above mentioned simple configuration, a storage queue or interface module priority, on the basis of which access to the common transmission medium is determined. In the prior art arrangement, these packet priorities are determined in function of the storage time of the packets in the storage queues and of at least one packet characterising parameter. Examples of such a packet characterising parameter are referenced in the prior art document in column 2, lines 4-9, as being for instance the packetising delay and/or the type of information contained in the analogous data packets filling the queue. The prior art arrangement however does not provide a solution for granting access to one of several identical storage queues, each receiving simultaneously a similar data input stream. This situation is for instance

present in access multiplexers, where in the upstream direction similar bitstreams carrying identical type of data packets and pertaining to identical quality of service categories, are competing for getting access on one common data bus. The prior art method and arrangement does not provide any solution for fairly granting access to one of these apparently similar bitstreams, arriving simultaneously at the different storage queues.

An object of the present invention is thus to provide a method and an arrangement for prioritised transmission of packets from a plurality of storage queues, each interfacing one corresponding input bitstream to one common transmission medium as described in the mentioned prior art document, but which provides a means for fairly determining access amongst similar input bitstreams, whereby this similarity is both related to content as well as to connection parameters associated to these bitstreams.

According to the invention, this object is achieved by the fact that the queue priorities on the basis of which access is determined, are dependent upon a respective weight associated to each corresponding bitstream, as is further stated in the characteristic part of claims 1 and 7.

In this way, by attributing a weight to apparently similar input bitstreams, entering at distinct storage queues, and by letting the priorities on the basis of which access is determined, to also be dependent upon these weights, a method for prioritised weighted access allowing to favour or discriminate some of these similar bitstreams with respect to the others, is realised.

An additional characteristic feature of the present invention is that the dependence of the queue priorities upon said weights is realised as is further described in claims 2 and 8.

In this way, by relating an amount of successive queue priorities with a weight, and by relating the access criterion with the sequence said successive priorities are sorted out, a bitstream with a relatively higher weight, will have a higher probability for accessing the transmission medium than a bitstream with a relatively lower weight, as will be explained more into detail in further paragraphs.

Another characteristic feature of the present invention is explained in claims 3 and 9.

Resetting the respective priorities at predetermined instances, results in priority cycles, during which, for each respective storage queue, the respective amount of priorities are compared, as will be explained more into detail in further paragraphs. Since the generation of each new priority is in fact related to the transmission of at least one packet from the associated storage queue, as was already known from the prior art arrangement, the subject method thus results in transmitting, for each storage queue, an amount of packets, directly related to its weight, during such a cycle. This means that the thus obtained method is not only simple, but guarantees, during each cycle, a fair access to the transmission

medium for each respective storage queue, in relation to the respective weight of the corresponding bitstream.

Yet another characteristic feature of the present invention is further described in claims 4 and 10.

In this way, a very simple relationship between the amount of values to select from said interval and said respective weights is obtained, thus resulting in a very simple and cheap arrangement.

Still a further characteristic feature of the present invention is described in claims 5 and 11.

Thereby, the general relationships for any predetermined value of the minimum and maximum limits of the weights and the resolutions of the priorities and the weights allow to determine appropriate values for said normalisation factor and for the limits of the interval.

Yet a further characteristic feature of the present invention is further described in claims 6 and 12

In this way, by letting the respective successive priorities associated to the corresponding bitstreams to be equidistant numbers spread over the entire width of the interval $[P_{min}, P_{max}]$, jitter is decreased as much as possible. Indeed, by smoothly distributing the priority values over the available range, a maximum variation amongst the different users is obtained.

The present invention also relates to a queue priority determining means, an access criterion and an access control device, such as those disclosed in the

above described new arrangement. The above mentioned and other objects and features of the invention will become more apparent and the invention itself will be best understood by referring to the following description of an embodiment taken in conjunction with the accompanying drawing which depicts an embodiment of an arrangement according to the invention.

An arrangement DTA for performing the subject method is depicted in the figure. This arrangement DTA is adapted to receive m input bitstreams, m being an integer not lower than 2. The bitstreams are respectively denoted BS1 to BSm, each respective one having an associated weight, denoted w_1 to w_m , carried by identically named weight input signals, being input signals to the arrangement. Data packets of each respective input bitstream BS1 to BSm are temporarily stored in respective storage queues denoted SQ1 to SQm, before being transmitted towards a common transmission medium, denoted UDATA. Each respective storage queue SQ1 to SQm is further associated to a respective queue priority determining means, denoted QPDM1 to QPDMm, included in the arrangement and adapted to generate a respective queue priority, QPR1 to QPRm. These queue priorities are used for determining which one of the m storage queues may transmit a data packet to the common transmission medium UDATA. This access arbitration is performed by an access control device, denoted ACD, of which various embodiments are known in the art. An embodiment of such an access control unit, as shown in the figure and similar to the one of the already referenced prior art document, consists of an

access interface, denoted AAM and being adapted to receive from each of the queue priority determining means QPDM1 to QPDMm the corresponding queue priorities QPR1 to QPRm, and to determine therefrom, in accordance to a predetermined access criterion, a winning priority denoted WP. This winning priority WP is then sent back towards m processing units, respectively denoted PR1 to PRm, included in the access control device ACD and associated to a respective storage queue SQ1 to SQm. These processing units are adapted to receive, besides the winning priority signal WP sent by the access interface AAM, also the respective priorities PR1 to PRm from the respective queue priority determining means QPDM1 to QPDMm. Furthermore, each respective processing unit is further adapted to compare both received input values, and to determine therefrom the value of a respective control signal, denoted G1 to Gm, being indicative of whether the associated respective storage queue SQ1 to SQm may transmit at least one data packet towards the common transmission medium UDATA. However other embodiments for determining the access arbitration based on the individual queue priorities exist, as is well known by a person skilled in the art. Since detailed realisations of such access control units are beyond the scope of this document, they will not be further described.

The invention concerns the method used for determining the respective queue priority values QPR1 to QPRm, in relationship to the predetermined access criterion. The method basically consists of determining, for each respective storage queue SQ1 to SQm, a respective amount, being a respective positive integer number, denoted K1 to Km, and determining this respective amount of successive queue priorities as a respective set of distinct values selected from a real interval $[P_{min}, P_{max}]$, and sorted out in a predetermined sequence. The predetermined access criterion is thereby related to this predetermined sequence.

The interval $[P_{min}, P_{max}]$ from which the respective queue priorities are selected, is identical for all storage queues. The limit values of this interval, P_{min} and P_{max} , are real numbers, with P_{min} being lower than P_{max} . These limit values of this interval are identical for all respective storage queues, and are in some embodiments determined by the user, during initialisation of the arrangement. In most embodiments however, P_{min} and P_{max} are design parameters, thereby imposing some absolute limits when choosing P_{min} and P_{max} . These absolute limits are generally used, since these allow for the largest flexibility for the weight values, as will become clear from further paragraphs. However selecting these limits is not mandatory. In either case, once P_{min} and P_{max} are determined, they are stored into a memory included in each queue priority determining means QPDM1 to QPDMm. More details about how to determine these values will be discussed in a further paragraph of this document.

In some embodiments, a distinct queue priority is determined and transmitted from the queue priority means towards the access control unit, each time a new packet is to be transmitted from the associated storage queue. This means that for bitstream Bm, the Km successively decreasing priority values, denoted here PRm1, PRm2, ..., PRmKm, are determined for Km successive packets to be transmitted from queue SQm. Each respective queue priority determining means QPDM1 to QPDMm is therefore informed by its respective storage queue SQ1 to SQm about a packet being sent, by means of respective control signal CS1 to CSm.

In other variant embodiments however, a queue priority may be associated to the transmission of a predetermined or programmable amount of packets to be transmitted. Thus each time one queue priority is selected as the winning priority, the associated storage queue may then transmit this predetermined amount of successive packets, instead on only one. Since this extension is easily realised by the addition of extra circuitry to the arrangement, which is well known to a person skilled in the art, and since this is not really relevant to the invention, the remainder of the document will only consider the case of one data packet being transmitted when the corresponding priority is the winning priority.

For each respective storage queue or bitstream BS1 to BSm, the respective amount K1 to Km of successively determined queue priority values, is related to the respective weights, w1 to wm, as input on the similarly named weight input lines. These weight values are real numbers, larger than 0. For any arbitrary bitstream BSi (not shown on the figure), the relationship between the respective amount Ki and the respective weight wi is very simple and consists of multiplying the respective weight wi with a normalisation factor NF, in accordance to formula (1):

$$K_i = w_i \cdot NF \quad (1)$$

This normalisation factor NF is dependent upon the resolution of the weights, denoted Rw, being a predetermined minimum distance between two consecutive weight values. This normalisation factor is further limited by a predetermined minimum limit value for the weights, denoted wmin. This relationship being formalised by:

$$NF \geq \max \{ 1/Rw, 1/wmin \} \quad (2)$$

whereby \geq indicates larger than or equal to,
1/Rw is the inverse of the predetermined resolution Rw of the weights,
1/wmin is the inverse of the predetermined minimum limit value wmin for the weights.

This formula (2) is obtained by considering that a minimum allowed difference in weight values, as

expressed by their resolution Rw, still has to result in a distinct amount of values, thus resulting in a minimum difference of 1. On the other hand, in the extreme case of a predetermined minimum limit value wmin for the weights, being lower than the predetermined resolution of the weights, even this bitstream should be able to get access, expressed by the fact that $NF \cdot wmin \geq 1$. By combining both requirements, formula (2) is obtained, stating that NF may not be lower than the maximum from $1/wmin$ and $1/Rw$.

The exact relationship between the interval width (Pmax-Pmin), the predetermined limit values of the weights, wmin and wmax, and the resolution of the priorities, denoted Rp and corresponding to the minimum distance between two consecutive priority values, is given by the following formula (3),

$$NF \cdot wmax \leq [(Pmax - Pmin) \cdot 1/Rp] + 1 \quad (3)$$

whereby \leq indicates lower than or equal to
1/Rp is the inverse value of the predetermined resolution Rp of the priorities,

this formula (3) reflecting that the maximum amount of distinct values to be selected from the interval [Pmin, Pmax] has to be lower than the number of priorities available within the interval [Pmin, Pmax].

From these formulas the boundary values for NF, Pmin, Pmax, and Rp can be determined, if the predetermined limit conditions of the weights are known on beforehand. The value of NF can be determined, for instance by means of a separate device included in the arrangement DTA. This device is adapted to receive the predetermined values of wmin and Rw, and to generate therefrom a value of NF, for transmitting this to a second device adapted to determine Pmin, Pmax and Rp from NF and wmax. In the opposite case of predetermined design limits for Pmin, Pmax and Rp, the boundary values for the weights may be derived, once a value for NF is selected, for instance by a person using this arrangement. In the arrangement depicted in the figure, the predetermined limit values of Pmin, Pmax and Rp are used, whereby a user has the freedom to select NF in accordance to formulas (2) and (3). In this case, Pmin, Pmax and Rp are design parameters stored in a memory location within or even inherent to each queue priority determining means, as will become clear from a further paragraph. Once the user has made a selection for NF, this value is programmed and stored within each queue priority determining means. Remark that the chosen value for NF, has to be applied for all queue priority determining means.

The principle for selecting the respective amount of priorities is illustrated in the following example, where 3 bitstreams BS1, BS2 and BS3, with associated weights w1 = 2, w2 = 3 and w3 = 5 are competing for access. The resolution for the weights Rw is 1 and is equal to the minimum weight limit wmin. The maximum weight

limit w_{\max} is 5. In this example, the limit values of the weights are thus known.

K1, K2 and K3 successive priorities attributed to each associated storage queue SQ1, SQ2, SQ3 are now to be selected from an interval $[P_{\min}, P_{\max}]$.

The normalisation factor NF has to be larger than or equal than 1 according to formula (2). Thus, in principle, any value not lower than 1 may be chosen as normalisation factor, either being a real or integer value. The choice of a value of NF, will however also determine some restrictions for determining the limit values of the interval and the resolution of the priorities, as was clear from formula (3). In this example NF will be chosen as 1. This means that respectively 2, 3 and 5 distinct values are to be selected from an interval, and arranged in a predetermined order. In this example a descending order is used.

Formula (3) gives the boundary restrictions for the determination of the interval. In fact a lot of different choices are possible. In case of a resolution of 1 for the priorities, the priorities are integer numbers and $P_{\max} - P_{\min}$ has to exceed 4. The interval $[1, 5]$ can then for instance be used as well as the interval $[10, 15]$ or even the interval $[100, 1000]$, since the restrictions only refer to the minimum width. In case of a smaller resolution of the priorities, for instance 0.1, even real values within the interval $[0, 1]$, or even negative real values within the interval

$[-4, -3]$, are possible, although for practical reasons negative values are rarely used.

If the resolution of the priorities is chosen to be 1, and if the interval $[1, 5]$ is chosen, 2 successive integer priorities for successive packets to be transmitted from SQ1 may be chosen as 5, 4 but nothing excludes to choose 2, 1 or 4, 2. In this example the values 2, 1 are chosen.

Similarly, for BS2, 3 successive integer priority values from the interval $[1, 5]$ may be 3, 2, 1 or 5, 3, 1 or any set of 3 distinct integer values, sorted in decreasing order, from the interval $[1, 5]$. The set 3, 2, 1 is chosen.

For bitstream BS3, there is only one way for selecting 5 distinct integer values, in descending order, from the interval $[1, 5]$, namely 5, 4, 3, 2, 1, since the interval boundaries correspond to the limit value expressed by formula (3).

The access criterion determining access to one of the storage queues based on their respective queue priorities, is related to the chosen sequence for ordering the distinct values from the interval, as this sequence represents successive respective queue priorities. In the example where the values were put in descending order, the access criterion basically consists of selecting the highest queue priority amongst all respective queue priorities as the winning priority, whereby the corresponding storage queue is granted the access. The winning priority WP is thus the highest amongst the respective priorities QPR1 to QPRm.

In case of two or more storage queues with equal

maximum priorities, a second selection has to be performed, this second selection being possibly related to a physical position of the storage queues or related slots on a board, as is mentioned in the published European Patent Application, publication number 0562222A1. Since this second selection is however outside the scope of this document, further implementation details will not be described in this document.

These sequences of respective queue priorities will then be further repeated at respective predetermined instances, related to the respective weights, as long as this respective weight has not changed. In general, since the value of these respective weights may vary in time, the respective queue priorities are reset at respective predetermined instances, to one of the values of the respective sequence, this respective sequence thereby corresponding to the actual value of the respective weight at this particular instance. Indeed, since the weight can vary in time, the number of values to be selected from the interval can vary as well, this resulting in a sequence being different from the previous one. And, even in case the respective weight has not changed, an alternative sequence of successive values may be chosen within the interval, at these predetermined instances. In general the first value of the new or actual sequence is selected for resetting to, but this is not always the case. For implementing this principle, several alternatives can be used. Two of them will be described, although they are not exhaustive.

A first alternative, performed by a first variant embodiment of the arrangement, consists of continuously repeating the first determined sequence of consecutive queue priorities, as long as the value of the respective weight does not change, and as long as the corresponding storage queue is not empty. Thus for an arbitrary storage queue SQi (not shown on the figure), successive consecutive groups of Ki queue priorities are continuously determined, resulting in the following succession of values: QPRi1, QPRi2, ..., QPRiKi, QPRi1, QPRi2, ..., QPRiKi, QPRi1, QPRi2, ..., QPRiKi etc. This means that for each new successive groups of Ki packets to be transmitted, the same priorities are applied in the same order, without interruption, as long as storage queue SQi is filled and as long as w_i is not changed.

In case of successive queue priorities put in descending order, the access criterion related to the first alternative method in this case not only consists of determining the winning priority WP to be the highest queue priority amongst all respective queue priorities, but also not exceeding an aim priority. Once the winning priority is determined, the aim priority is changed to the value of the thus obtained winning priority. When it is detected by the access control device ACD that all respective queue priorities are higher than the aim priority, the latter is reset to P_{\max} .

The initial value of this aim priority may be set at P_{\max} , but this is not mandatory.

In the previous example, applying this first alternative method will result in the following queue priorities :

BS1 : 2,1,2,1,2,1,2,1,2,1,....etc.

BS2: 3,2,1,3,2,1,3,2,1,3,2,1,....etc.

BS3: 5,4,3,2,1,5,4,3,2,1,5,4,....etc.

The related access criterion will then result in the following access sequence :

- first BS3 gets the access because of its highest priority value (5), the aim is determined to be 5.
- next BS3 again wins the access, because of its highest priority (4), the aim is changed to 4,
- next both BS2 and BS3 have the highest priority, in this case a further discrimination has to take place by a second access criterion as was previously described. This may for instance be based on the position, as was previously described as well. Suppose that for this second access criterion BS3 then gets priority over BS2 which gets priority over BS1.

BS3 thus wins. The aim becomes 3.

- next BS2 is served, the aim is not changed.
- next BS1, BS2, BS3 all have the winning priority (2). BS3 thus gets the access, the aim is changed to 2.
- next BS2 gets access, aim is not changed,
- next BS1 gets access, aim still not changed,
- next BS3 is granted access. The aim is changed to 1
- BS2 and BS1 are then consecutively granted access, despite the higher priority value of BS3. However this value (5) exceeds the aim.

After BS1 has won the last access, the next competing priorities are 2,3, and 5, all definitely higher than the last updated aim value of 1. From that moment on, the aim priority is reset to 5 by the access control device, and the same procedure is repeated.

As a matter of fact, this first alternative method has introduced a sequence of priority cycles, in time separated by the resetting of the aim priority to Pmax, where this time is determined by the sum of all respective amounts, thus also the sum of all respective weights. Since the values of the weights may vary in time, and since new bitstreams may become activated or added, the duration of each of these cycles may thus vary as well.

From this example, one can also conclude that in the first priority cycle SQ1 has sent 2 packets, SQ2 has sent 3 packets and SQ3 has sent 5 packets, which is exactly corresponding to a distribution of the bandwidth in accordance to the weights, during this cycle. Thus a fair and weighted access arbitration scheme is established during each priority cycle by this method.

It is also observed that during this cycle BS3 has sent 3 consecutive packets, this resulting in a rather

bursty transmission of data. In order to reduce this, the respective priority values are selected as equidistant values spread over the entire width of the interval [Pmin,Pmax]. For any arbitrary storage queue SQi, this distance may be obtained by dividing the interval width (Pmax-Pmin) by (Ki-1), and attributing respective queue priorities to SQi starting from Pmax, $P_{\max} - (P_{\max} - P_{\min}) / (K_i - 1)$, $P_{\max} - 2(P_{\max} - P_{\min}) / (K_i - 1)$, ..., and ending with $P_{\min} + (P_{\max} - P_{\min}) / (K_i - 1)$, Pmin, and continuously repeating the same sequence as long as the value of the weight wi does not change. Another possibility consists of selecting the following sequence : $P_{\max} - (P_{\max} - P_{\min}) / 2K_i$, $P_{\max} - 3(P_{\max} - P_{\min}) / 2K_i$, ..., $P_{\min} + 3(P_{\max} - P_{\min}) / 2K_i$, $P_{\min} + (P_{\max} - P_{\min}) / 2K_i$.

In the first variant embodiment, adapted for performing this first alternative method, the respective queue priority determining means QPDM1 to QPDMm include a device adapted to determine the respective amounts K1 to Km from the normalisation factor NF and from the respective weights, input to these respective queue priority determining means. Each respective queue priority determining means further includes a selection device for selecting, or calculating in case of equidistant values, the thus determined respective amount of values from the predetermined interval of which the limits were stored in a memory included within each queue priority determining means. Each queue priority determining means further includes means for sorting out the selected values in the appropriate sequence, and the necessary control logic for outputting a next priority value from the thus obtained sequence, upon a packet being sent from the associated storage queue. They further include means for resetting their respective queue priority back to the first one of the sequence upon detecting that a cell was sent while the queue priority was the last one of the sequence. In case of equidistant priorities, the queue priority determining means may be simpler, including a rotating counter, counting down, starting at for instance Pmax, with an decrement of $(P_{\max} - P_{\min}) / (K_i - 1)$, and upon detecting that a packet was sent while the priority was at its minimum value Pmin, starting again. In this case, Pmax and Pmin are thus inherent design parameters of the counters of the queue priority means themselves. In these embodiments, the queue priority determining means further include a device for calculating the respective decrement values of their counters, from the value of NF and the respective weights.

In case of the second mentioned sequence for having equidistant values, the limit values of the counters need to be adaptable as well, resulting in more complex embodiments.

The access control unit adapted for performing this first variant of the method needs to include the necessary devices adapted to determine the aim priority, and to further adapt it, in accordance with the previously described first variant of the method. Since a person

skilled in the art is further adapted to realise detailed implementations from the above described functional description of the access control device and the queue priority determining means, further details will not be described in this document. It however needs to be remarked that in all possible implementations, necessary control logic is to be foreseen for synchronising the different signals present in the arrangement. Again, this kind of circuitry is commonly known to a person skilled in the art and will therefore not be further described in this document.

A second alternative, realised by a second variant of the embodiment, consists of also continuously repeating the first determined sequence, while however inserting an auxiliary priority Paux, between two consecutive sequences. Again, this repetition is only performed on the condition that the value of the respective weights does not change and that the associated storage queue not gets empty. Thus for an arbitrary storage queue SQi (not shown on the figure), basically successive consecutive groups of $K_i + 1$ queue priorities are continuously generated, resulting in the following succession of values :
 QPRi1, QPRi2, ..., QPRiKi, Paux, QPRi1, QPRi2, ..., QPRiKi, Paux, etc.

In case of successive queue priorities QPRi1 to QPRiKi put in descending order, the access criterion related to the second alternative method merely consists of selecting the highest queue priority amongst all respective queue priorities, whereby however in case the winning priority WP is determined to be Paux, all storage queues are inhibited from transmitting a packet, and each queue priority determining means has to reset its respective queue priority to the first one of its respective sequence. To this purpose, all queue priority determining means are informed upon the fact that $WP = Paux$, by means of a separate control signal (not shown on the figure), generated by the access control device ACD.

The auxiliary priority Paux has to lie outside the interval $[Pmin, Pmax]$; in the case of successive priority values sorted out in descending order, Paux has to be smaller than Pmin.

In the previous example, applying this second alternative method, whereby Paux is arbitrary chosen to be 0.5, the following queue priorities result :

BS1 : 2, 1, 0.5, 2, 1, 0.5, 2, 1, 0.5, etc
 BS2 : 3, 2, 1, 0.5, 3, 2, 1, 0.5, 3, 2, 1, 0.5 etc.
 BS3 : 5, 4, 3, 2, 1, 0.5, 5, 4, 3, 2, 1, 0.5, etc

The access criterion related to this second alternative method will then result in the following access sequence : BS3, BS3, BS3, BS2, BS3, BS2, BS1, BS3, BS2, BS1. During the next comparison period, all three queue priorities will have the value 0.5, which is now determined to be the winning priority WP. The access control device thereby inhibits all storage queues from

transmitting a data packet and generates an extra control signal for transmission to all queue priority determining means, which then, upon receiving this signal, reset themselves for generating the first value of the original sequence etc.

Again, this second alternative method has resulted in a sequence of priority cycles, during which successive values of the K_1, \dots, K_m priorities are compared. In this case, the onset of such a priority cycle is not determined by the value of an aim priority, but by the value of the winning priority WP, which, upon being equal to Paux, triggers the simultaneous resetting of the individual storage queue priorities itself. Again, one may notice that during this priority cycle, BS3 has transmitted 5 packets, BS2 has transmitted 3 packets and BS1 has transmitted 2 packets. This principle again results in a fair distribution of the fairness, based on the values of the weights at that moment.

The queue priority determining means adapted for performing this second alternative of the method is basically similar to the one performing the first alternative method, but includes additional means adapted to insert the auxiliary priority value after each determined set of values within the interval $[Pmin, Pmax]$, and other additional means for resetting itself to the first value of the sequence upon receiving the extra control signal from the access control device indicating that the winning priority is Paux. This access control device of this second variant embodiment does not include means for determining the aim priority, as was the case for the first variant embodiment, but includes the necessary circuitry for generating control signals to all queue priority determining means upon detecting that WP is Paux, and means for inhibiting all storage queues from transmitting a data packets upon detecting the same.

Other, more sophisticated extensions or deviations from this second alternative method, realised by yet other variant embodiments, are possible as well. One of them provides the possibility that, even in the case that $WP = Paux$, one of the storage queues having a queue priority being Paux, may transmit a packet, after which step all queue priorities are reset to the first value of their respective sequences, except for the storage queue just having transmitted the packet. This particular one is then immediately reset to the second one of its range.

In all variants of the second alternative method, the duration of a priority cycle is still again determined by the sum of the respective amounts of successive priorities from the interval $[Pmin, Pmax]$, which is related to the sum of the respective weights at the onset of this priority cycle. The exact duration of these cycles is further determined by implementation details of the access control device, queue priority determining means, and the storage queues themselves, which are all synchronised by synchronisation logic included in the arrangement (not shown on the figure).

This second alternative method in principle also

allows to choose at the onset of each new cycle, a different set of values from the interval, or even to update the interval. The same could however be realised with the first alternative method, at the start of a new cycle. Since these extensions however involves extra complex circuitry, and since they do not bring any advantage to the invention, these extensions are hardly used.

As for the first alternative, burstiness is seriously reduced by taking equidistant values for the respective successive priorities determined within one cycle. The same values as previously mentioned for the first alternative may be used.

For both alternatives, the amount of packets being transmitted during one priority cycle from one particular storage queue SQ_i , is thus directly proportional to the weight w_i of the corresponding bitstream. The ratio between the share of bandwidth between 2 individual bitstreams BS_i and BS_j (not shown on the figure), is therefore also directly proportional to the ratio between their respective weights w_i and w_j . Fairness is thus obtained. In case of more complex embodiments where more than one packet is transmitted from a storage queue with the winning priority, it is mandatory for guaranteeing fairness that the amount of packets transmitted from each queue that won the access, is the same, when using the subject method.

The subject method is especially valuable for dividing the remaining bandwidth, according to a specific weight, amongst the active users of the Unspecified Bit Rate, abbreviated by UBR, service category. This kind of service category is standardised by the ATM forum and intended for non-real time applications such as file-transfer and electronic mail and is in fact a best-effort service, for which only the residual bandwidth, remaining after other, well-specified quality of service categories are served. Nevertheless, it may be the wish of the telecom operator to yet control the partition of this available bandwidth among the active users of this UBR category, thereby favouring some users with respect to others, for instance based on a different tariffing rate. In the latter case, the users prepared to pay more are offered a relative larger share of the remaining bandwidth. As a matter of fact, these difference in tariffing schemes results in a difference in relative weights the operator has attributed to these users for these connections.

In general, these weights on the basis of which an operator may wish to discriminate amongst several user are not necessarily related to the tariffing, but may in other situations may as well be related to other characteristics of the users such as their profile (business, domestic etc.) or to their fidelity. These reasons however are outside the scope of the invention, and will therefore not be further discussed.

It further needs to be remarked that the queue priority generating means only output a value of the determined sequence for non-empty queues. In case of an empty queue, a predetermined empty-queue priority,

having a predetermined value being lower than P_{min} or even P_{aux} for the second alternative method, can then for instance be used. The necessary circuitry for generating this is rather trivial and will therefore not be further described. All described variants of the method, as well as all described embodiments of the arrangement, easily allow for new connections or bitstreams, with associated new weights, to be added or activated, while yet preserving fairness. Indeed, as long as the weights of the newly added or activated bitstream lie within the predetermined weight limits $[w_{min}, w_{max}]$, a corresponding amount of queue priorities may be determined according to formulas (1), (2) and (3). For the second alternative, new or newly activated connections will start with priorities at P_{aux} , for synchronising reasons. Whenever a new cycle is started, their priorities will be automatically be reset to the first value of the sequence, according to the second alternative method. The delay that is encountered is only low and can be afforded. With the first alternative, when a new bitstream becomes activated, in order to preserve fairness, transmitting its determined sequence of priority values has to be postponed to the onset of the next priority cycle, by extra circuitry to be added to for instance the queue priority determining means.

It further needs to be remarked that, in stead of selecting the respective amount of queue priorities from the interval $[P_{min}, P_{max}]$ and putting them in descending order, they may as well be put in ascending order. In the latter case, the access criterion of course has to be changed such that the lowest amongst all priorities is selected for determining the winning priority. All variants of the method, described for the case of descending priorities, may now as well be applied to the case of ascending priorities, whereby some slight changes have to be introduced, e.g. the aim priority needs to be reset to P_{min} in the first variant of the method, and in the second variant of the method the auxiliary priority value P_{aux} now has to be higher than P_{max} . The necessary changes for adapting the different embodiments to the case of ascending priority values are rather trivial and will therefore not be further described into detail.

While the principles of the invention have been described above in connection with specific apparatus, it is to be clearly understood that this description is made only by way of example and not as a limitation on the scope of the invention.

Claims

1. Method for prioritised transmission of packets from a plurality of storage queues (SQ_1, \dots, SQ_m), each interfacing one corresponding input bitstream (BS_1, \dots, BS_m) to one common transmission medium (UDATA), said method including a step of determining a respective queue priority (QPR_1, \dots, QPR_m) for each respective storage queue (SQ_1, \dots, SQ_m), on the basis of which access to said

common transmission medium (UDATA) is determined in accordance to a predetermined access criterion,

characterised in that

said respective queue priority (QPR1,...,QPRm) is dependent upon a respective weight (w1,...,wm), being a positive real number and associated to said corresponding input bitstream (BS1,...,BSm).

2. Method according to claim 1

characterised in that

the dependency of said respective queue priority (QPR1,...,QPRm) upon said respective weight (w1,...,wm) relates to a respective amount (K1,...,Km), of successive values of said respective queue priorities, said respective amount of successive values being a respective set of said respective amount of distinct values, selected within an interval [Pmin,Pmax] and sorted out in a predetermined sequence, with Pmin being a first real number lower than a second real number Pmax, said predetermined access criterion thereby being related to said predetermined sequence.

3. Method according to claim 2

characterised in that

said respective queue priority is reset, at respective predetermined instances related to said respective weights, to a respective value at a predetermined position of said respective set.

4. Method according to claim 2

characterised in that

said respective amount (K1,...,Km) is determined as the integer most approximating the product of said respective weight (w1,...,wm) and a normalisation factor (NF).

5. Method according to claim 4

characterised in that

said normalisation factor is determined as a value not lower than the maximum between the inverse value (1/Rw) of a predetermined resolution (Rw) of said weights and the inverse value (1/wmin) of a predetermined minimum limit value (wmin) of said weights, whereby 1 added to the product of the width of said interval (Pmax-Pmin) and the inverse value (1/Rp) of a predetermined resolution (Rp) of said priorities is not lower than the product of a predetermined maximum limit value (wmax) for said respective weights and said normalisation factor (NF).

6. Method according to claim 2

characterised in that

for said respective amount (K1,...,Km) being larger than 1, said respective amount of distinct values are equidistantly spaced over the total width of

said interval [Pmin,Pmax].

7. Arrangement (DTA) for prioritised data transmission of packets from a first plurality of storage queues (SQ1,...,SQm) included in said arrangement, each interfacing one corresponding input bitstream (BS1,...,BSm) to one common transmission medium (UDATA) coupled to said plurality of storage queues, said arrangement (DTA) including a second plurality of queue priority determining means (QPDM1,...,QPDMm), each respective queue priority determining means (QPDM1,...,QPDMm) of said second plurality being associated to a respective storage queue (SQ1,...,SQm) of said first plurality and adapted to determine a respective queue priority (QPR1,...,QPRm), said arrangement (DTA) further including an access control device (ACD) adapted to receive from said each respective queue priority determining means (QPDM1,...,QPDMm) said respective queue priority (QPR1,...,QPRm) and to determine therefrom, in accordance to a predetermined access criterion, which storage queue of said first plurality gets access to said common transmission medium (UDATA),

characterised in that

said each respective queue priority determining means (QPDM1,...,QPDMm) is further adapted to determine said respective queue priority (QPR1,...,QPRm) from a respective weight (w1,...,wm), being a positive real number and associated to said corresponding input bitstream (BS1,...,BSm).

8. Arrangement according to claim 7

characterised in that

said each respective queue priority determining means (QPDM1,...,QPDMm) is further adapted to determine a respective amount (K1,...,Km) of successive values of said respective queue priorities as a respective set of said respective amount of distinct values, selected within an interval [Pmin,Pmax], and sorted out in a predetermined sequence, said respective amount (K1,...,Km) thereby being dependent upon said respective weight (w1,...,wm) associated to said corresponding input bitstream (BS1,...,BSm), and Pmin being a first real number smaller than Pmax being a second real number, whereby said predetermined access criterion is related to said predetermined sequence.

9. Arrangement according to claim 8

characterised in that

said each respective queue priority determining means (QPDM1,...,QPDMm) is further adapted to reset said respective queue priority at respective predetermined instances related to said respective weights, to a respective value at a predetermined

position, of said respective set.

10. Arrangement according to claim 8

characterised in that

said each respective queue priority determining means (QPDM1,...,QPDMm) is further adapted to determine said respective amount (K1,...,Km) as the integer most approximating the product of said respective weight (w1,...,wm) and a normalisation factor (NF). 5 10

11. Arrangement according to claim 10


characterised in that

said arrangement further includes a device adapted to determine said normalisation factor (NF) as a value not lower than the maximum between the inverse value (1/Rw) of a predetermined resolution (Rw) of said weights and the inverse value (1/wmin) of a predetermined minimum limit value (wmin) of said weights, said arrangement further including a second device adapted to determine said first real number Pmin and said second real number Pmax so that 1 added to the product of the width of said interval (Pmax-Pmin) and the inverse value (1/Rp) of a predetermined resolution (Rp) of said priorities is not lower than the product of a predetermined maximum limit value (wmax) for said respective weights, and said normalisation factor (NF). 15 20 25

12. Arrangement according to claim 8

characterised in that

said each respective queue priority determining means (QPDM1,...,QPDMm) is further adapted to determine said respective amount of successive values of said respective queue priorities as equidistantly spaced real values over the total width of said interval [Pmin,Pmax], in case said respective amount is larger than 1. 30 35 40 45 50 55





European Patent
Office

EUROPEAN SEARCH REPORT

Application Number
EP 97 40 1433

DOCUMENTS CONSIDERED TO BE RELEVANT			
Category	Citation of document with indication, where appropriate, of relevant passages	Relevant to claim	CLASSIFICATION OF THE APPLICATION (Int.Cl.6)
X	CHANGHWAN OH ET AL: "PRIORITY CONTROL ATM FOR SWITCHING SYSTEMS" IEICE TRANSACTIONS ON COMMUNICATIONS, vol. E75 - 8, no. 9, 1 September 1992, pages 894-904, XP000321325 * paragraph 3 *	1,2,6-8, 12	H04L12/56 H04L29/06
X	US 5 218 676 A (BEN-AYED MONDHER ET AL) * abstract * * column 10, line 26 - line 56 * * figure 9A *	1,7	
A	HLUCHYJ M G ET AL: "QUEUEING DISCIPLINES FOR INTEGRATED FAST PACKET NETWORKS" DISCOVERING A NEW WORLD OF COMMUNICATIONS, CHICAGO, JUNE 14 - 18, 1992, vol. 2 OF 4, 14 June 1992, INSTITUTE OF ELECTRICAL AND ELECTRONICS ENGINEERS, pages 990-996, XP000326820 * paragraph 3.2 * * paragraph 3.3 *	1,7	
A	ANONYMOUS: "Weighted Queueing Algorithm for Efficient Asynchronous Transfer Mode Traffic Shaping" IBM TECHNICAL DISCLOSURE BULLETIN, vol. 39, no. 4, April 1996, NEW YORK, US, pages 161-164, XP002048400 * the whole document *	1,4,5,7, 10,11	H04L
D,A	EP 0 584 405 A (ALCATEL NV ;BELL TELEPHONE MFG (BE)) * the whole document *	1-12	
The present search report has been drawn up for all claims			
Place of search THE HAGUE		Date of completion of the search 27 November 1997	Examiner Perez Perez, J
<p>CATEGORY OF CITED DOCUMENTS</p> <p>X : particularly relevant if taken alone Y : particularly relevant if combined with another document of the same category A : technological background O : non-written disclosure P : intermediate document</p> <p>T : theory or principle underlying the invention E : earlier patent document, but published on, or after the filing date D : document cited in the application L : document cited for other reasons & : member of the same patent family, corresponding document</p>			

EPO FORM 1503 03 82 (P04C01)